

INTERNAL USE ONLY ADDRESSES

BACKGROUND OF THE INVENTION

5 Network equipment manufacturers assign 48-bit identifiers called media access control (MAC) addresses to network devices they make. MAC addresses are allocated to manufacturers in blocks by the Institute of Electrical and Electronics Engineers, Inc. (IEEE) Registration Authority. All MAC addresses in an allocated block include the same three-byte "organizationally unique identifier" (OUI) reserved for the manufacturer.

10 The allocation of MAC addresses with OUIs is intended to encourage interoperability of network devices. If network devices made by different manufacturers applied in the same network are to communicate unambiguously through exchange of MAC-addressed frames, each network device must be assigned a MAC address that is unique. And since it is not known *a priori* in which network a device will be applied, each manufacturer must assign to a network device a MAC address which is unique among not only its own manufactured devices, but unique among network devices made by other manufacturers as well. Such global uniqueness is guaranteed by allocating MAC addresses having OUIs.

15 Although most MAC addresses are used for communication between network devices, as outlined above, use of MAC addresses exclusively within a single network device is becoming more common. One such "internal only" use of MAC addresses arises in the context of exceptional frame forwarding within a LAN switch. Generally speaking, a LAN switch is a multi-port LAN device which

20

25

interconnects LAN end-stations residing in different LAN broadcast domains, or LAN segments, through different ports on the switch. Such a switch typically forwards conventional frames received from a source end-station over a first port to a destination end-station over a second port based on MAC addressing information encoded in the frame and the switch's knowledge of the port through which the destination end-station can be reached. Such a switch may, however, depart from convention when forwarding certain exceptional frames. To perfect exceptional frame forwarding, MAC addresses may be temporarily assigned to interfaces or ports of the switch and encoded in exceptional frames upon internal transmission to cause such frames to be captured by such interfaces or ports. Such temporarily assigned and encoded MAC addresses are never "seen" on an external LAN transmission medium; they are only "seen" internally.

Because MAC addresses applied in "internal only" uses are never "seen" outside the network device to which they are assigned, maintaining their global uniqueness among network equipment manufacturers to avoid addressing ambiguities is not necessary. Manufacturers nonetheless often assign globally unique MAC addresses (i.e. allocated by the IEEE Registration Authority and having the manufacturer's globally unique OUI) for "internal use only" applications in network devices, reducing the number of globally unique MAC addresses remaining in the available MAC address pool unnecessarily.

As described in Internet Engineering Task Force Requests for Comment (IETF RFCs) 1062, 1122 and 1700, it is known in the context of Internet Protocol (IP) addressing to reserve the network number "127" for a host "loopback"

function for use on any IP host such that any IP datagram having a network "127" address loops-back inside the host and never appears on any IP network. However, IP host loopback addressing is directed to a Layer 3 (i.e. routing) protocol. Moreover, there is no suggestion to assign loopback addresses to
5 switching devices or applying such addresses in special processing of inbound datagrams.

As described in RFC 2464, it is also known to locally administer MAC addresses by setting a "U/L" bit at a particular bit position in the 48-bit address. However, locally administered MAC addresses are not reserved for "internal
10 only" uses and may be encoded in frames transmitted on external LAN media.

Accordingly, it would be desirable to implement a reservation, allocation, assignment and application protocol for networking addresses, such as MAC addresses, which takes advantage of the fact that addresses used exclusively in internal applications would not cause addressing ambiguities even if assigned to
15 a plurality of devices in the same network.

SUMMARY OF THE INVENTION

In one aspect, the present invention provides reservation, request, allocation, assignment and application protocols for conserving networking addresses in a finite address domain wherein a portion of the addresses must be
20 uniquely allocated. In accordance with the protocols, a first set of addresses within the domain are reserved for "internal use only" and are allocable to a plurality of organizations for assignment and application in a plurality of devices without risking the introduction of addressing ambiguities into communication

networks. A second set of addresses within the domain and not having the "internal use only" restriction are reserved for a single organization and are assigned and applied in a single network device to avoid risking the introduction of addressing ambiguities into networks. The "internal use only" addresses are preferably "intra-switch only" addresses and are allocable to a plurality of organizations for assignment and application in a plurality of switches. The addresses are preferably MAC addresses.

These and other aspects of the present invention may be better understood by reference to the following detailed description taken in conjunction with the drawings briefly described below.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a conceptual level diagram illustrating request, allocation assignment and application protocols for MAC addresses in accordance with a preferred embodiment of the invention;

Figure 2 is a block diagram of a LAN switch in which MAC addresses allocated in accordance with a preferred embodiment of the invention may be assigned and applied;

Figure 3 is a diagram of a conventional frame upon receipt by the LAN switch according to Figure 2 from a LAN in an exemplary application protocol;

Figure 4 is a diagram of a conventional frame upon transmission on the backplane of the LAN switch according to Figure 2 in an exemplary application protocol;

Figure 5 is a diagram of a conventional frame upon transmission by the LAN switch according to Figure 2 to a LAN in an exemplary application protocol;

Figure 6 is a diagram of an exceptional frame upon receipt by the LAN switch according to Figure 2 from a LAN in an exemplary application protocol;

5 Figure 7 is a diagram of an exceptional frame upon transmission on the backplane of the LAN switch according to Figure 2 by a network interface in an exemplary application protocol;

Figure 8 is a diagram of an exceptional frame upon retransmission on the backplane of the LAN switch according to Figure 2 by the management interface;

10 Figure 9 is a diagram of an exceptional frame upon transmission by the LAN switch according to Figure 2 to a LAN in an exemplary application protocol;

Figure 10 is a flow diagram of reservation, request and allocation protocols for MAC addresses in accordance with a preferred embodiment of the invention; and

15 Figure 11 is a flow diagram of assignment and application protocols for MAC addresses in accordance with a preferred embodiment of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

In Figure 1, preferred reservation, request, allocation and application protocols for MAC addresses are shown at a conceptual level. MAC address domain 100 includes address ranges reserved for LAN equipment manufacturer LEM1 110 (LEM1R addresses), for LAN equipment manufacturer LEM2 120 (LEM2R addresses) and for "intra-switch only" applications (ISR addresses). LEM1R and LEMR2 address ranges are reserved by assigning LEM1 110 and

LEM2 120 a three-byte identifier (OUI) and following a convention that all addresses having an OUI as the first three bytes are reserved for the LAN equipment manufacturer to which the OUI has been assigned. Thus, for example, all MAC addresses in the range "00-20-DA-xx-xx-xx" may be reserved for Alcatel Internetworking, Inc., the assignee hereof, by assigning Alcatel Internetworking, Inc. the OUI "00-20-DA". ISR addresses may be used by any LAN equipment manufacturer but are reserved for "internal use only" within LAN switches. ISR addresses may be reserved by assigning an N-byte identifier outside the range of any OUI assigned to a LAN equipment manufacturer and following a convention that all addresses having the N-byte identifier are reserved for "internal use only".

In a preferred request-and-allocation protocol, LAN equipment manufacturers request MAC address allocation and are, generally speaking, allocated blocks of MAC addresses within the range of their assigned OUIs. Thus, as shown, LEM1 110 requests allocation of MAC addresses and in response is allocated a block of addresses within the LEM1R address range. LEM2 120 requests allocation of MAC addresses and in response is allocated a block of addresses within the LEM2R address range. However, it is an important feature of the invention to depart from this basic request-and-allocation protocol with respect to MAC addresses for application exclusively within a LAN switch. For such applications, rather than allocating blocks of addresses within a LAN equipment manufacturer's dedicated address range, addresses from the shared ISR address range are used.

In a preferred assignment protocol, LAN equipment manufacturers assign MAC addresses to LAN switches. Thus, as shown, LEM1 110 assigns LEM1R addresses to LEM1 switch 115 which may be applied internally and externally. Similarly, LEM2 120 assigns LEM2R addresses to LEM2 switch 125 which may be applied internally and externally. However, in an important feature of the invention, LEM1 110 and LEM2 125 also assign ISR addresses to their respective switches 115 and 125 solely for internal application. The ISR addresses assigned by LEM1 110 and LEM2 120 to their respective switches 115 and 125 are from the shared ISR address range within MAC address domain 100 and, therefore, may (but do not necessarily) include one or more identical addresses.

In a preferred application protocol, the MAC addresses assigned to LAN switches are applied in a network. Thus, as shown, LEM1R and LEM2R addresses are transmitted between and within switches 115 and 125 and LAN 135 in MAC-addressed frames, whereas ISR address transmission is restricted to within switches 115 and 125, respectively. Since ISR addresses are not transmitted by switches 115 and 125 on the external LAN medium (interconnecting switches 115 and 125 via LAN 135), there is no risk that any ISR address assigned to both switch 115 and switch 125 will cause addressing ambiguity in the network. Yet by assigning and applying potentially redundant ISR addresses in switches 115 and 125 for "intra-switch only" applications, unique addresses from domain 100 are advantageously conserved. LAN switches 115 and 125 are multi-port devices having as one functional element

the ability to make forwarding decisions as a function of MAC addresses, and may fall within any of various categories of data switching elements including hubs, bridges, switches, routers and gateways.

Various elaborations of the reservation, request, allocation, assignment and application protocols heretofore described are possible, as described hereinafter. Nevertheless, at a fundamental level, these basic protocols, despite their apparent simplicity, are believed to confer significant advances over the prior art.

Referring now to Figure 2, a LAN switch in which MAC addresses allocated in accordance with a preferred embodiment of the invention may be assigned and applied is shown. Switch 200 includes a matrix of frame buses 211-219 driven by interfaces 201-209, respectively. Interfaces 201-209 include network interfaces 201-208 connected to respective LANs 221-228 via respective transmission media 231-238, and management interface 209. Each bus has a root interfacing with the one of interfaces 201-209 having the exclusive right to transmit frame data on the bus (i.e. the root interface) and leaves interfacing with the plurality of interfaces 201-209 receiving frame data off the bus (i.e. the leaf interfaces). Preferably, each interface is the root interface on one of buses 211-219 and is a leaf interface on all buses 211-219, including the bus for which it is the root interface. Buses 211-219 are broadcast-oriented such that all data bursts propagated on a bus reach all interfaces 201-209. In addition to transmitting and receiving frame data, management interface 209 serves as the "nerve center" of switch 200 which assists network interfaces 201-208 in special frame processing,

including processing of Spanning Tree frames as hereinafter described. Of course, the root-to-leaf architecture described above is one of many possible architectures for a switch operative in accordance with the present invention. Other architectures may have a single common bus between interfaces or "full mesh" matrix providing point-to-point connections between all combinations of interfaces.

Interfaces 201-209 perform filtering checks on frames received of buses 211-219 by looking at MAC addresses in a local header appended to frames. Frames not having as a MAC address an address known by interfaces 201-209 are dropped, or "filtered", subject to certain exceptions. One exception applies when none of interfaces 201-209 recognizes the MAC address. Interfaces 201-209 share the results of filtering checks to avoid filtering frames whose MAC address has not been learned by any interface. Such "unknown destination" frames are captured by all interfaces. More particularly, in an exemplary filtering check, an interface applies the following filtering rules:

Rule 1: If the frame has a MAC address known on the interface, the filtering check is passed. The frame is captured.

Rule 2: If the frame has a MAC address not known on the interface, and the MAC address is known on another interface, the filtering check is failed. The frame is filtered.

Rule 3: If the frame has a MAC address not known on the interface, and the MAC address not known on another interface, the filtering check is passed. The frame is captured.

5 A frame in this context may be any protocol data unit having a MAC header with a destination MAC address therein. A frame may in certain circumstances encapsulate an entire Network layer protocol data unit (e.g. an entire IP packet). By way of example, frames may be Ethernet protocol data units and Spanning Tree protocol data units.

10 An exemplary application of conventional and exceptional frame forwarding will now be described by reference to frames traversing switch 200. In Figures 3 through 5, an exemplary application of conventional frame forwarding, which does not apply an ISR address, is illustrated. Frame 300, originated by an end-station on a LAN associated with one of network interfaces 201-208 and destined for an end-station on a LAN associated with a different one of network interfaces 201-208 arrives at switch 200. Frame 300 includes a MAC header, an 802.1Q tag, an IP header and additional information not relevant to the present example. The MAC header includes, among other information not shown, a destination MAC address of the destination end-station; in this case a globally unique address allocated to LAN equipment manufacturer X (LEMXR₁).
15 The network interface receiving frame 300 from the LAN appends a local header to frame 300 which includes the MAC address LEMXR₁ retrieved from the MAC header, resulting in frame 400. Frame 400 is propagated on the one of buses

211-218 for which the receiving network interface is the root interface. Frame 400 is received from the one of buses 211-218 by all interfaces 201-209, which perform filtering checks individually by looking at the MAC address in the local header. The MAC address $LEMXR_1$ is known on the network interface through which the destination end-station is connected to switch 200 (generally as a result of a "source learning" operation performed on a previous frame transmitted by the destination end-station), therefore filtering Rule 1 (above) applies and the frame is captured by such interface. The MAC address $LEMXR_1$ is not known on the other network interfaces or management interface 209, therefore filtering Rule 2 (above) applies and the frame is filtered by such interfaces. The capturing network interface strips the local header from the frame, resulting in frame 500, which is forwarded out the capturing network interface to the destination end-station.

In Figures 6 through 9, an exemplary application of exceptional frame forwarding, which applies an ISR address ISR_1 , is illustrated. In the exemplary application, the exceptional forwarding enables dynamic path selection. Frame 600 includes a MAC header and additional information not relevant to the present example. The MAC header includes, among other information, a destination MAC address; in this case a dynamic path selection address (DPS_1). The network interface receiving frame 600 applies a local header to frame 600 which includes the MAC address DPS_1 retrieved from the MAC header resulting in frame 700. Frame 700 is propagated on the one of buses 211-218 for which the receiving network interface is the root interface. Frame 700 is received from the

one of buses 211-218 by all interfaces 201-209, which perform individual filtering checks by looking at the MAC address in the local header. The MAC address DPS_1 is known on management interface 209, which has been configured to recognize the address. Therefore filtering Rule 1 applies and frame 700 is captured by management interface 209. The MAC address DPS_1 is not known on network interfaces 201-208, however. Therefore filtering Rule 2 applies and frame 700 is filtered by network interfaces 201-208. Management interface 209 resolves, based on the current state of a dynamic path selection algorithm, on which one of network interfaces 201-208 the frame must be forwarded by switch 200 and generates a frame 800 having a local header including "intra-switch only" MAC address ISR_1 , which the resolved forwarding network interface has been preconfigured to recognize. The destination MAC address in the MAC header remains DPS_1 . Frame 800 is propagated on the backplane, this time on bus 219 for which management interface 209 is the root interface. Frame 800 is received by all interfaces 201-209, which perform individual filtering checks by looking at the MAC address in the local header. The MAC address ISR_1 is known on the resolved forwarding network interface; therefore, filtering Rule 1 applies and the frame is captured by such interface. The MAC address ISR_1 is not known on the other network interfaces or management interface 209; therefore filtering Rule 2 applies and the frame is filtered by such interfaces. The resolved network interface strips the local header from frame 800 resulting in frame 900, which is forwarded out the resolved network interface. In the foregoing manner, an "intra-switch only" MAC address is advantageously applied to effectuate dynamic path

selection, resulting in conservation of a globally unique MAC address which would otherwise be used. It will be appreciated that each network interface may be configured to recognize a different "intra-switch only" MAC address such that multiple globally unique MAC addresses are conserved.

5 Turning now to Figure 10, a flow diagram illustrating preferred reservation, request and allocation protocols for MAC addresses is presented. Separate and distinct MAC address ranges are reserved for LAN equipment manufacturer X (LEMXR) and for "intra-switch only" applications (ISR) (1010). A request for allocation of MAC addresses from a LAN equipment manufacturer X is received
10 (1020), in response to which the LAN equipment manufacturer is allocated a block of addresses from its reserved address range (LEMXR addresses) (1030).

Turning finally to Figure 11, a flow diagram illustrating preferred of assignment and application protocols for MAC addresses is presented. LAN equipment manufacturer X assigns MAC addresses to a LAN switch, including (i)
15 for internal and external application, allocated MAC addresses reserved for LAN equipment manufacturer X (LEMXR addresses), and (ii) for "internal use only", MAC addresses reserved for "intra-switch only" applications (ISR addresses) (1110). The switch is interconnected in a network (1120) and booted-up (1130). In network operation, LEMXR addresses are applied internally and externally,
20 whereas ISR addresses are applied exclusively within the switch (1140).

It will be appreciated by those of ordinary skill in the art that the invention can be embodied in other specific forms without departing from the spirit or essential character hereof. The present description is therefore considered in all

respects illustrative and not restrictive. The scope of the invention is indicated by the appended claims, and all changes that come within the meaning and range of equivalents thereof are intended to be embraced therein.

002740 29805560